

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<small>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</small> PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 06-30-2009		2. REPORT TYPE Final		3. DATES COVERED (From - To) 01-01-2006 - 12-31-2008	
4. TITLE AND SUBTITLE Active Computation for Networked ATR Systems				5a. CONTRACT NUMBER N00014-06-1-0061	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Joseph A. O'Sullivan Michael D. DeVore				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Washington University				8. PERFORMING ORGANIZATION REPORT NUMBER 22-1331-59621	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 875 North Randolph Street Arlington, VA 22203-1995				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>Our work focused on design and implementations guidelines for ATR systems that must adapt to time-varying resource constraints. The goal is to have systems that can dynamically change based on the availability of time, number of processors, communication bandwidth, and system architecture including access to remote databases. The systems should have nearly optimal performance given the resources available. One set of designs proposed is based on hierarchical data and processing models and on-line performance evaluation. The hierarchical models are based on information-theoretic considerations, a fundamental basis that is unique to our approach.</p>					
<div style="display: flex; justify-content: space-between;"> <div>15. SUBJECT TERMS</div> <div style="font-size: 2em; font-weight: bold;">20090701221</div> </div>					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON Joseph A. O'Sullivan
a. REPORT U	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code) 314-935-4173



Washington University in St. Louis

SCHOOL OF ENGINEERING & APPLIED SCIENCE

Department of Electrical and Systems Engineering

Joseph A. O'Sullivan
Samuel C. Sachs Professor

June 30, 2009

Grant Number: N00014-06-1-0061

Title: **Active Computation for Networked ATR Systems**

PI: Joseph A. O'Sullivan, Washington University

Subcontract: Michael D. DeVore, University of Virginia and Barron Associates, Inc.

Final Report

Abstract

Our work focused on design and implementations guidelines for ATR systems that must adapt to time-varying resource constraints. The goal is to have systems that can dynamically change based on the availability of time, number of processors, communication bandwidth, and system architecture including access to remote databases. The systems should have nearly optimal performance given the resources available. One set of designs proposed is based on hierarchical data and processing models and on-line performance evaluation. The hierarchical models are based on information-theoretic considerations, a fundamental basis that is unique to our approach.

Introduction

In this project, we developed a fundamental basis for the design and analysis of automatic target recognition (ATR) systems based on information-theoretic considerations. The theory developed could provide a new way of analyzing ATR systems. In the view of the principal investigator, this theory provides a fundamental starting point for determining what is possible in ATR system design. The algorithms are based in this information-theoretic analysis, and indicate that the most informative parts of models should be considered first, along with the most informative aspects of the measurements. Our theory points to methods for measuring that information and performing the design.

In practice, the theory needs to be adapted to real world considerations. In that direction, there need to be methods for on-line performance prediction. We have developed, as part of this work, new techniques that are practical (computationally fast) for predicting this performance.

We have explored the application of the methods in multiple scenarios, including point-cloud data and optical imaging systems, in addition to radar imaging sensors that we have previously considered.

The ATR systems we envision work in dynamic environments, where the system on which data must be processed changes over time. That is, there may be different sensors, databases may have dynamic availability or bandwidth connecting them to the rest of the system. Sensors may be on mobile platforms that have dynamic communication links.

The algorithms have a goal to account for evolving system architecture, time constraints, processor availability, and communication bandwidths to achieve nearly optimal performance under a wide range of operating conditions. A core idea is to process the most informative target features followed by successively less informative ones. Operational requirements and environment conditions dictate the extent of feature processing, so processing continues until an answer is required from the system. Better performance (lower error rate) can be achieved as more detail is used, but additional resources are consumed (resources such as time, bandwidth, and processing).

In a previous report, we noted that:

Both design and operation guidelines for systems employing successively refinable algorithms can be developed by relating the accuracy of a system to the resources consumed while rendering a decision. From a design perspective, a framework is provided through component modeling for comparing the resulting performance of very different architectures, such as one with very little processing on-board the sensor platform but relatively high communication bandwidth, versus one with more sophisticated computations on-board, but with relatively low communication bandwidth. This type of modeling enables the system designer to quantitatively understand the implications of design choices that are meaningful at the application level. The models relate the accuracy of a recognition system, the time required to solve a recognition problem, and the resources that must be brought to bear. From an operational perspective, the analysis results can be used to help allocate resources appropriately.

Comments on Funding

This project was originally funded at a level of \$500,000 over three years. After a modification in the funding, a total of \$268,735 was allocated. This led to a reduced scope of the project relative to the proposed effort. Nevertheless, we feel that we have made fundamental and lasting contributions to the theory and practice of the design of automatic target recognition systems. Our concept that ATR systems be designed to adapt on the fly to the resources available is revolutionary. Properly transferred to operational ATR systems, this idea could lead to dramatic improvements in ATR system performance in the field. In this final report, we outline the essence of our results, and we describe publications that resulted from this support.

Key Theoretical Contribution: Information-Theoretic Bounds for Recognition Systems

There are several theories that enable the design and development of automatic target recognition systems that adapt to available resources. One enabling theory is to bound the achievable performance of any ATR system subject to constraints on resources. A second is to bound the performance of any ATR system that can be achieved successively, that is by achieving one performance level using some resources and then achieving a better level of performance using incrementally more resources. We have made significant progress on both of these theories. Elsewhere in this report, we discuss significant advances that we have made in computing

performance on the fly in order to be able to approximate the performance achieved at any given point.

ATR systems always operate under resources constraints. These constraints may include any or all of the following:

- *Constraints on the amount of time available for performing an ATR function.* Most ATR systems are deployed in time-critical situations. There may be threat or potential threat in the environment. There is a critical need for fast and accurate ATR. However, not every deployment scenario has the same constraint on time. One scenario may require an answer in half the time that another demands. The design of ATR systems flexible enough to adapt to varying amounts of time is one of the goals of this project. The characterization of potential ATR performance as a function of time is a canonical situation explored for quantifying the tradeoff between resources and performance.
- *Constraints on the size of the database used to store data models.* Given an object, there is either a model or a set of templates (either is referred to as a model) used within an ATR system. The model should include the variability of the data as a function of target aspect angle, position, and articulation relative to the sensor. The complexity of that model determines the size of the database required. If the model is more complex, then the database size required to store it is greater, and typically the time required to assess whether the data fit that model is greater. Also, if there are a fixed number of targets, available data storage must be divided among them. If the model is too complex, then there may not be enough time for an ATR system to be able to use the entire model. This motivates the need for the model to be stored in a flexible way so that the most informative aspects of the model can be accessed first, then the less informative aspects next. In this context, informative is measured in terms of impact on ATR performance.
- *Constraints on the bandwidth of the communication link connecting the sensor to the ATR processor.* Often the processor performing ATR computations and the sensor that collects the data used for ATR are not collocated. There is a need to transmit the data from the sensor to the ATR processor. This communication link may have finite bandwidth that limits the performance. If the sensor is on board an unmanned aerial vehicle (UAV), and is operating in a hostile environment, then the communication link may be jammed as well. Thus the data rate may fluctuate and may be unpredictable in advance. In these scenarios, it is important for the ATR system to be able to operate with whatever data are provided. In addition, the sensor platform would need to transmit the most informative bits first. The most informative bits in this context are those that impact the performance of the ATR system the most.
- *Constraints on the number of targets that can be explored within the ATR system.* ATR systems are often designed to consider a finite number of targets. As the number of targets grows, the overall system may become more complex. A central issue that our theory contributes to is the tradeoff between the number of targets or target classes and the achievable performance. We note here, that the number of potential target scenarios that can be distinguished is exponentially large in the number of degrees of freedom of the data measured. These target scenarios include target classes, targets within a class, and configurations (position, orientation, and articulation) of a given target.
- *Constraints on the number of processors available for performing ATR computations.* If the processing is performed on board a UAV, the availability of processors may be

dynamic. There may be more critical needs that require processing (such as survival). Much of the processing in ATR can be done in parallel, so the availability of additional processors directly increases the performance achieved in a given amount of time. In that sense, the number of processors has an impact on performance directly analogous to the impact of time.

- *Other resource constraints that are functionally equivalent to one of the above.*

The research question that we addressed related to the first theoretical contribution is quantifying bounds on achievable performance given constraints on resources. We believe that we have made fundamental contributions to a fundamental version of this problem. Our paper with Brandon Westover that appeared in the January 2008 IEEE Transactions on Information Theory describes these results.

M. Brandon Westover and Joseph A. O'Sullivan, "Achievable rates for pattern recognition," *IEEE Transactions on Information Theory*, Volume 54, Issue 1, Jan. 2008 pp. 299 – 320.

See also

M.B. Westover and J.A. O'Sullivan, "Pattern Recognition on a Bit Budget," to appear in *Neural Information Processing Systems*, N. Tishby, editor, 2009.

The successive refinement aspects have been submitted for publication and are now in revision. A preliminary version of the successive refinement results are available in

J.A. O'Sullivan, N. Singla, and M.B. Westover, "Successive Refinement for Pattern Recognition," *2006 IEEE Information Theory Workshop*, Punta del Este, Uruguay, pp. 141-145, March 13-17, 2006.

A linear pattern matching approach was explored in

Po-Hsiang Lai and Joseph A. O'Sullivan, "Pattern Recognition System Design with Linear Encoding for Discrete Patterns," *2007 IEEE International Symposium on Information Theory*, Nice, France

Related results in biometric performance prediction were in

N.A. Schmid and J.A. O'Sullivan, "Performance Prediction Methodology for Multi-Biometric Systems," in Springer-Verlag book on biometrics, A. Abidi and B. Abidi editors, 2006 (in press).

Transitions to DoD, Government, and Industry:

- We interacted with Alan Van Nevel throughout the project, discussing the needs of networked ATR systems in the context of systems of interest for the Navy, especially those of Navair in China Lake.

- We met with industry contacts on numerous occasions to discuss the potential impact of our work in realistic deployment scenarios. Feedback from these contacts helped guide the directions of our work. Companies include Boeing, ESSI, and Lockheed Martin. In addition, through our attendance at conferences, we had many more informal interactions with industry.
- We met with representatives of the Department of Homeland Security to discuss the potential impact of our work in homeland security. The Center for Security Technologies at Washington University provided a conduit for access to many key contacts in security areas. Our work was presented at annual reviews of the Center for Security Technologies, attended by representatives throughout the government and industry.

Publications and Selected Presentations:

We gave at least two presentations at ONR workshops and meetings

Joseph A. O'Sullivan gave an invited presentation at an Army Research Office meeting on automatic target recognition held at the University of Michigan in July 2007

Our paper that appeared in the October 2008 issue of IEEE Trans. Aerospace and Electronic Systems is: "Analysis of data and model accuracy requirements for target classification using lidar data," by Xin Zhou and Michael D. DeVore, *IEEE Transactions on Aerospace and Electronic Systems*, Volume 44, Issue 4, Oct. 2008, pages 1416 – 1432. It takes a detailed look at sensing requirements of a lidar-based ATR system in several practical deployment scenarios.

The research leading to the PhD dissertation of Xin Zhou was funded primarily from this ONR project. The title is "Statistical Model-Based Object Recognition from Three-Dimensional Point-Cloud Data," Dissertation for Doctor of Philosophy in Systems Engineering, University of Virginia, May 2008. Most of his work was in the area of target recognition from lidar data.

The work in the MS thesis of Shankha Basu. Basu was primarily funded through this ONR grant. The title is "A High-Level Modeling Framework for Total System Performance of Automated Target Recognition in Distributed Surveillance Systems," Master of Science in Systems and Information Engineering, University of Virginia, May 2008. He focused on questions of distributed computation of ATR likelihoods on miniaturized computing hardware. He looked at recognition performance as a function of computational constraints, like number and speed of processors, etc. This work was based on gray-scale imagery simulated from 3D CAD models using the rendering engine POV Ray.

The paper "Statistical Models for Target Detection in Infrared Imagery," by Samuel H. Huddleston, Xin Zhou, William B. Evans, Alice Chan, and Michael D. DeVore, was presented at the SPIE symposium in 2007. It addresses target detection in infrared imagery. More accurately, it illustrates a philosophy for approaching ATR problems. This paper grew out of a class project Michael D. DeVore assigned in his fall 2006 pattern recognition course. That assignment was motivated by an academic challenge problem formulated by the ATR technical group of SPIE. As a result, it has kind of an introductory feel that is not typical of most journal papers. The ONR project funded Xin Zhou's time on the paper as well as that of Michael DeVore. The lead author (Sam Huddleston) is an Army Major who was assigned to the systems engineering department at

the University of Virginia to earn his MS. He is now assigned to teaching duties at West Point. The other two students working on this paper were unfunded.

A previous paper was presented in 2006 at the SPIE symposium. It forms the basis for some of the work on this project. It addresses minimum probability of error recognition from ladar data in the presence of significant clutter. This work is also included in the PhD dissertation of Xin Zhou. "Minimum Probability of Error Recognition of Three-Dimensional Laser-Scanned Targets," by Michael D. DeVore and Xin Zhou, *Proc. SPIE Automatic Target Recognition*, Vol. 6234, 2006.

Elements of robustness analysis to take into account modeling inaccuracies are explored in our paper.

Lichun Li and Joseph A. O'Sullivan, "Robustness via a Tradeoff Between Fisher Information and Relative Entropy," 2007 IEEE/SP 14th Workshop on Statistical Signal Processing, 26-29 Aug. 2007 pp. 239 – 243.

More details are available in the 2007 dissertation of Lichun Li for the doctor of science degree in electrical engineering at Washington University.